

Chapter 1

Introduction

1.1 Motivation

The research presented in this thesis is primarily motivated by the desire to provide a *framework* for adaptive hypertexts which can establish user interests based on paths of traversal, rather than using pre-stated interests alone. This motivation also leads to the requirement that the hypertext be adaptable so that its community of users can participate in the creation of an environment which can then automatically assist individual users in their search for relevant information.

When we use large scale hypertext systems, such as the World Wide Web, to browse for information, we are faced with a decision-making problem each time we access a document which does not completely satisfy our information need. The problem posed is concerned with what we should do next to locate more relevant information. If the document has links to other documents, which, if any, should we follow? Should we revisit a previously accessed document and follow a different link from there? Are we able to construct a query which accurately represents the information we seek, so we can submit that query to a search engine? Is there better information available than that which we have already located? These problems, together with navigational difficulties such as how to return to a previously visited document, and how the different documents that we accessed are organisationally related to each other, are largely the result of a lack of link semantics, a proliferation of authors with different authoring styles, and mismatches between what an information provider and an information consumer consider relevant. On the other hand, this lack of rigidity and formality are among the most appealing aspects of the World Wide Web.

In this thesis, we attempt to solve these problems pragmatically without imposing standards such as link semantics and an ontology. The approach we take is that the community of users of a hypertextual information system (the information consumers) actively contribute to the re-organisation and description of existing information according to how they consume it. Users can then be automatically and individually assisted to describe their information need and to guide them to relevant information.

1.2 HyperContext

HyperContext is a new framework for adaptable and adaptive hypertext which actively supports users while they browse and search for information. Our basic premise is that the same piece of information can have different, and sometimes even contradictory, subjective meanings to different readers, and that these meanings can be at least partially described¹. We call these different descriptions *interpretations* of information.

In a hypertext, the same document can occur in many locations as the destination of many links. We capture the description of the document at each location as a different interpretation of the same document. Whenever a document is accessed, it is first interpreted prior to being displayed to the browser. The influencing factor that causes a document to be interpreted is called a *context*. Different interpretations of the same document can present the user with different views of the hyperspace, with each interpretation of a document potentially containing different links with different destinations. Interpreting a document causes the hyperspace to be *partitioned* into the set of (probably) relevant documents and (probably) non-relevant documents. Partitions are generated as a user browses through hyperspace, and we call the traversed portion of a partition a *context session*.

When a link is traversed during a context session the link destination is interpreted prior to being presented. The interpretation consists of out-links and their destinations and a description of the document in the current context. The same document may have a different description and set of out-links and link destinations for each context in which it exists.

HyperContext uses the different possible interpretations of a document to ascertain a user's interests. HyperContext can also guide the user to relevant information by recommending which link to follow in each document on a path.

¹ A very simple exercise which demonstrates this premise is for the reader to consider what this chapter means *before* and *after* reading the rest of the thesis.

HyperContext is a social hypertext - users can create new interpretations of information, and modify existing ones. Users can also create new links, as well as adding new documents to the hyperspace.

1.3 Thesis overview

1.3.1 Chapter 2

This chapter identifies some of the problems associated with navigating through a hypertext which we aim to tackle and alleviate through the HyperContext framework. A number of metaphors are introduced which will guide us through the earlier chapters of the thesis.

Hypertexts are implicitly social in nature, but the community of users normally cannot directly influence the organisation of information to make it easier for others to locate the information they seek. When the same document occurs in different locations in the hyperspace it is normally re-used *in situ*, so different browsers who access the same document having followed different paths to it are presented with the same set of links. Good interface design dictates that users should not be overloaded with choice, so document authors must *a priori* decide on the number and destinations of links included in their contribution. Different browsers may find different subsets of those links relevant, while others may find them all non-relevant. Traditional hypertexts generally cannot directly provide any assistance to the user to locate the information they seek, except through information retrieval systems which require users to accurately describe their need, even though the user may not be able to do so.

HyperContext is briefly introduced as a framework for adaptable and adaptive hypertext which allows the community of users to directly modify the structure and organisation of information in the hyperspace, and which provides individualised adaptive navigation support to its users.

1.3.2 Chapter 3

We place the work described in this thesis into its relevant research area, which crosses the boundaries of hypertext, information indexing and retrieval, and user modelling. We give a brief description of typical general-purpose hypertext and ubiquitous information retrieval systems, comparing their structure, representation, organisation of information, and the levels of interaction the user has with each. We then focus specifically on adaptive hypertext systems. Brusilovsky ([15], [14]) has written recent and excellent literature

reviews for Adaptive Hypertext Systems (AHSs), which we use as our starting point. We notice that AHSs used in education (Intelligent Tutoring Systems) and domain-independent WWW-based AHSs are based on different underlying assumptions, and the approaches they take to achieve adaptivity are substantially different. ITS-based AHSs can model the user and the domain very closely, while heterogeneous WWW-based AHSs tend to be unable to do so. Additionally, while ITS-based AHS users can have an explicitly stated goal (and a plan which may be available to help achieve that goal), the WWW-based AHS user often has an unstated or unclear goal, but adaptive assistance needs to be given to the user within a few link traversals. We suggest that there is a difference between a short-term user interest and a long-term interest, and that the type of support an AHS should provide is quite different in each case, primarily because a user with a short-term interest is unlikely to be capable of describing her information need as well as she would a long-term interest.

1.3.3 Chapter 4

We describe the structure of the HyperContext framework to support adaptability and adaptivity. We support adaptive hypertext through the use of multiple *interpretations* of information. An interpretation is a description of a document in *context*. We maintain that different interpretations of the same document support different descriptions of the same document according to its location in the hyperspace. We associate links with interpretations, rather than documents. Each interpretation of the same document can have different links, and even if they have the same links, the links can have different destinations in each interpretation. Interpretations of documents in context allow a hyperspace to be partitioned into documents which are relevant (because they can be reached in the current context) and those which are non-relevant (and cannot be reached). A user session can be composed of different *context sessions* - sessions in which links are traversed in context. We automatically detect the end of a context session, even though the user session may not have ended.

The HyperContext framework has three layers. The *Object Layer* contains document descriptions to bind HyperContext document references to the documents stored in the universe outside HyperContext. The *Structure Layer* represents the interpretations of documents and the links between them, and the *Presentation Layer* supports interaction between the user and HyperContext.

1.3.4 Chapter 5

Chapter 5 describes the HyperContext framework from the user's point of view. First, we describe how a hypertext can be adapted by the user adding document descriptions to the Object Layer, and creating and modifying document interpretations and links in the Structure Layer. Then we describe how the hypertext adapts to a user while she browses. During a context session, a model of the user's short-term interests is constructed. From the user model, a query can be automatically formulated on the user's behalf. The HyperContext framework utilises the services of an external information retrieval system, through an interface. Following a search request, the results can be processed in several ways. The HyperContext framework supports three search paradigms: *Traditional Information Retrieval* searches throughout the hypertext for information relevant to a user query; *Information Retrieval-in-Context* limits its search to a *context sphere* of interpretations reachable from a particular document interpretation; and, *Adaptive Information Discovery* automatically generates a query from the user model and searches throughout the hypertext. If a relevant document is found within the context sphere, the user can be guided to it by following recommended links. We construct the short-term user model by combining *salient interpretations* of documents accessed during the context session using a modification to the Rocchio method of relevance feedback.

1.3.5 Chapter 6

We consider the treatment of context in HyperContext. We give a brief literature review of logics of context and the applications of context, with particular reference to John McCarthy's research at Stanford University. Context is the environment within which information is understood. We show that this supports our thesis that information can have different interpretations in different contexts, and we identify the environment in hypertext that can provide context. We theorise about the possibility of automatically interpreting information in arbitrary contexts, and we decide that although HyperContext's Adaptive Information Discovery can provide dynamic links to interpreted information we prefer HyperContext's community of users to create the interpretations. We distinguish between *contextual relevance* and *superficial relevance*. Interpretations are contextually relevant minimally when they co-exist on the same context path, and they are superficially relevant if a matching algorithm (such as the cosine similarity measure) considers them relevant but they exist on different context paths. We discuss how the salient interpretation of a document, introduced in Chapter 5, is derived in terms of context, and introduce a weighted scale of confidence for incorporating salient interpretations into the short-term user model.

1.3.6 Chapter 7

We describe a client-server approach to a prototype of the HyperContext framework to demonstrate the adaptive support that can be given using interpretations of documents in context. We also show how HyperContext can be reduced to the equivalent of a non-adaptive hypertext such as the WWW.

1.3.7 Chapter 8

Chapter 8 describes the process of converting a Web site into a HyperContext hypertext. We show how an extended Boolean information retrieval system provides HyperContext with external information retrieval services. We use the resulting hypertext to demonstrate the advantages of a HyperContext hypertext over a non-adaptive hypertext. The experiments and results are described in Chapter 9.

1.3.8 Chapter 9

We describe three experiments and their results. The first experiment determines the method for deriving salient interpretations and combining them to form a short-term user model. The short-term user model is then used to select previously unseen relevant information. The second experiment shows that the salient interpretation can be used to locate relevant information faster than a single, re-usable description of information. We compare the results of brute-force search in adaptive and non-adaptive versions of a HyperContext hypertext and then again with the same searches employing a link ordering technique based on salient interpretations. The third experiment asks users at the end of a context session to give relevance judgements for documents recommended by HyperContext's Adaptive Information Discovery and a control method.

1.3.9 Chapter 10

Chapter 10 summarises the major work and findings of the thesis, and suggests further work on the HyperContext framework and the prototype. We also identify some other research opportunities.

1.4 Summary of contributions

It is generally accepted that information can be interpreted differently in different contexts. To the best of my knowledge, the explicit use of multiple interpretations of the same documents in a hypertext and associating links with these interpretations are novel.

Previous attempts have been made to determine a user's interests based on paths of traversal. I have taken the Rocchio method [74] to automatically re-formulate an information retrieval query based on relevance feedback and modified it derive a salient interpretation of a document in context. I believe that this particular extension to existing work, automatically deriving a salient interpretation to describe the likely user interest in the current document, and the combination of salient interpretations of documents accessed during a context session to construct a model of the user's short-term interests are novel.

The idea of adaptive hypertext is not new. However, the HyperContext framework as a whole and the manner in which adaptive support in HyperContext is provided are original. Early work on HyperContext has previously been published in [84] and [85], although this thesis contains substantial modifications to the approach.

SWISH-E [48], an extended Boolean Information Retrieval system, has been used to provide the HyperContext prototype with information retrieval services. SWISH-E is property of the Hewlett-Packard Company, who distribute the source code with no warranty. Permission is granted for the free use of the software, and modifications to the source code, under the terms of the GNU Public License [41]. The software modifications we have made are not for re-distribution and SWISH-E does not form part of the HyperContext framework or the prototype.